

Stochastic Flips on Dimer Tilings

Thomas Fernique and Damien Regnault

LIF, CNRS & Univ. de Provence
39 rue Joliot-Curie 13453 Marseille – France
{thomas.fernique,damien.regnault}@ens-lyon.org

Abstract. This paper introduces a Markov process inspired by the problem of quasicrystal growth. It acts over dimer tilings of the triangular grid by randomly performing local transformations, called *flips*, which do not increase the number of identical adjacent tiles (this number can be thought as the tiling energy). Fixed-points of such a process play the role of quasicrystals. We are here interested in the worst-case expected number of flips to converge towards a fixed-point. Numerical experiments suggest a $\Theta(n^2)$ bound, where n is the number of tiles of the tiling. We prove a $O(n^{2.5})$ upper bound and discuss the gap between this bound and the previous one. We also briefly discuss the average-case.

Introduction

Tilings are often used as a toy model for quasicrystals, with minimal energy tilings being characterized by local properties called *matching rules*. In this context, a challenging problem is to provide a theory for quasicrystals growth. One of the proposed theories relies on a relaxation process ([5] p. 356): a tiling with many mismatches is progressively corrected by local transformations called *flips*. Ideally, the tiling eventually satisfies all the matching rules and thus shows a quasicrystalline structure. It is compatible with experiments, where quasicrystals are obtained from a hot melt by a slow cooling during which flips really occur (Bridgman-Stockbarger method). It is however unclear whether only flips can explain successful coolings or if other mechanisms should be taken into account. This is deeply related with the convergence rate of such a flip-correcting process.

A cooling process aiming to be physically realist is described in [1]. It considers so-called *cut and project tilings* of any dimension and codimension, and performs flips which modify by ΔE the energy of the tiling with a probability proportional to $\exp(-\Delta E/T)$, so that the stationary distribution at fixed temperature T is the Boltzmann one.

A simplified cooling process is obtained by performing equiprobably at random only flips whose corresponding ΔE is above a fixed threshold. It has been studied on tilings of dimension one and codimension one (two-letter words) in [2]. We here focus on tilings of dimension two and codimension one (dimer tilings).

The paper is organized as follows. Sec. 1 introduce notations and basic definitions. We then describe, in Sec. 2, the cooling process we consider. We also formally state the convergence time we want to bound, and make conjectures based on numerical experiments. Sec. 3 is then devoted to the proof of an upper bound on the convergence time. This proof relies on a concentration result for some well-chosen function. We conclude the paper by a short section discussing prospects of this work, namely the non-tightness of the obtained theoretical bound, and the average convergence time, *i.e.*, when the initial tiling is chosen at random (instead of considering the one with the greatest convergence time).

1 Settings

Dimer tiling

Let $(\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3)$ be the unit vectors of the Euclidean plane of direction $\frac{\pi}{6} + k\frac{2\pi}{3}$, $k = 1, 2, 3$. They generate the so-called triangular grid. A *dimer* is a lozenge tile made of two adjacent triangles of the grid. A *domain* is a connected subset of the grid. Then, a *dimer tiling* is a tiling of a domain by dimers (Fig. 1, left).

Lift

Let $(\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3)$ be the canonical basis of the Euclidean space. The *lift* of a dimer tiling is its image by a map ϕ which is linear over tiles and satisfies $\phi(\mathbf{x} + \mathbf{v}_k) = \phi(\mathbf{x}) + \mathbf{e}_k$ for any two vertices \mathbf{x} and $\mathbf{x} + \mathbf{v}_k$ connected by an edge of the tiling. This map is uniquely defined up to a translation; we assume that $(0, 0, 0)$ belongs to the image of the domain boundary. Tiles are thus mapped onto facets of three-dimensional unit cubes, and dimer tilings are mapped onto stepped surfaces of \mathbb{R}^3 . This can be easily seen by shading tiles (Fig. 1, center).

Height

Following [7], we define the *height* of a point in a dimer tiling by the distance of its image under ϕ to the plane $x + y + z = 0$. We also define the *height* of a tile as the height of its center: this yields a third representation of dimer tilings, where tile colors depend on the height (Fig. 1, right).




Fig. 1. Three representations of the same dimer tiling.

Errors and energy

An *error* in a dimer tiling is an edge shared by two identical tiles (up to a translation). Error-free regions thus perfectly alternate tiles. The corresponding subsets of the triangular grid can be tiled by six-triangle hexagons, and the lift of an error-free region approximate the plane $x + y + z = 0$. In particular, all the tiles have the same height. The *energy* $E(\omega)$ of a dimer tiling ω is its total number of errors.

Islands and holes

Suppose that a dimer tiling ω contains a finite, connected and simply connected set of tiles σ whose boundary edges are errors and whose domain can be tiled by six-triangle hexagons. The number $A(\sigma)$ of such hexagons is called the *area* of σ . One checks that the height of tiles with an edge on the boundary of σ take only two values, say $h(\sigma)$ and $h(\bar{\sigma})$, respectively depending whether the tile is in σ or not. One says that σ is an *island* if $0 < h(\bar{\sigma}) < h(\sigma)$, a *hole* otherwise. In both cases, $h(\sigma)$ is called the *height* of σ . Fig. 2 illustrates this.




Fig. 2. Boundary tiles of this dimer tiling have the same height, say 0. From left to right: an island of height 1 and area 10; the same island with a hole of height 0 and area 2; once again the same island with a superposed island of height 2 and area 4; an island of height -1 and area 10 (which looks like a hole because its height is negative).

Volume

The *volume* $V(\omega)$ of a dimer tiling ω is defined as the sum of areas of its islands minus the sum of areas of its holes. Error-free tilings thus have volume zero. One checks that if the domain is finite, simply connected and admits an error-free tiling, then the errors of any dimer tiling of this domain are the boundaries of its islands and holes. Hence, tilings of volume zero are error-free. Otherwise, non-closed paths of errors can run across the domain, with endpoints (if any) on the boundary of the domain.

Partial order

A dimer tiling ω is said to be smaller than or equal to a dimer tiling ω' of the same domain, written $\omega \leq \omega'$, if each point has, in modulus, a height in ω smaller than or equal to its height in ω' . The set of tilings of a domain becomes a distributive lattice, whose extremal elements are the tilings of extremal volume.

Flip

Whenever a vertex x of a dimer tiling belongs to exactly three tiles, one gets a new dimer tiling by translating each of these tiles along the edges shared by the two other ones (or, equivalently, by rotating these tiles by $\frac{\pi}{3}$ around x). This local rearrangement of tiles is called a *flip* in x (Fig. 3). In the lift, it corresponds to add or remove one cube. One checks that a flip modifies the volume of a tiling by $\Delta V = \pm 1$ and its energy by $\Delta E \in \{0, \pm 2, \pm 4, \pm 6\}$. The total number of flips which can be performed on a tiling ω is denoted by $F(\omega)$.




Fig. 3. A flip and all its possible immediate neighborhoods (up to a rotation), with the according variation of energy of the tiling being reported. The four leftmost cases correspond to merging of island or creation of holes. In the cooling process considered in Sec. 2, only flips which do not increase the energy will be allowed.

Note that the volume of a tiling turns out to be the minimal number of flips required to transform this tiling into an error-free tiling (if the domain admits such a tiling).

Flip-accessibility

In [7], Thurston has shown that any two dimer tilings can be connected by performing a sequence of flips, *i.e.* are mutually *flip-accessible*, if they tile the same finite and simply connected subset of the triangular grid. The case of dimer tilings of the whole triangular grid has been considered in [3]. Here, the physical motivations exposed in the introduction lead us to focus on a *constrained* flip-accessibility, where only flips which do not increase the number of errors are allowed. Moreover, we are not interested in a flip-accessibility where flips can be carefully chosen, but in the case where flips are *randomly* performed. The next section makes this more precise.

2 The cooling process

Let us fix a finite, connected and simply connected subset of the triangular grid which does admit an error-free dimer tiling made of n tiles. Let Ω_n denotes the set of all the possible dimer tilings of this subset (flip-accessibility yields that all these tilings have n tiles).

The so-called *cooling process* we are interested in is the Markov chain defined over Ω_n as follows. It starts from $\omega_0 \in \Omega_n$ and produces a sequence $(\omega_t)_{t=1,2,\dots}$ where ω_{t+1} is obtained by performing on ω_t a flip, uniformly chosen at random among the flips which do not increase the number of errors¹. If there is no such flip, then the process stops (the tiling is said to be *frozen*).

Fig. 4 shows the cooling of a tiling of a hexagonal-shaped domain of 3779 tiles, chosen uniformly at random among the tilings of this domain. In this particular case, the process stops in 4290 steps, with the obtained frozen tiling turning out to be error-free. How general is such an evolution?




Fig. 4. Snapshots each 1000 flips of the cooling of a random tiling (from left to right).

The following proposition will give us a first insight into this evolution:

Proposition 1. *Whenever a tiling in Ω_n has an error, a flip which do not increase the number of errors can be performed onto. Moreover, this flip can be chosen so that it decreases the volume of the tiling.*

Proof. The existence of a (finite) error-free tiling ensures that errors form the boundary of islands. Consider an island whose height is, in modulus, maximal. Such an island is made of hexagons tiled by three-tiles. Performing a flip on such a hexagon decreases the area of the island (hence the volume of the tiling) and decreases the number of errors by $2(k-3)$, where k is the number of error edges of this hexagon. It thus suffices to find a hexagon with at least three error edges. Let us follow clockwise the boundary of the island: the direction between two consecutive edges changes by $\pm \frac{\pi}{6}$, and there is more negative variations because the boundary is closed. We can thus find two consecutive negative variations: the three corresponding edges are errors on the boundary of the wanted hexagon. \square

This proposition ensures that a frozen tiling is necessarily error-free. Moreover, it shows that the process almost surely stops in a finite time, since the

¹ In other words, the ΔE threshold discussed in the introduction is equal to zero.

volume of tilings in Ω_n is uniformly bounded and decreases with probability at least $\frac{1}{n}$ at each step.

In order to describe more precisely this cooling process, we are interested in the probability distribution of the random variable T , called the *convergence time*, which counts the number of steps before it stops. Here, we focus only on the *worst expected convergence time* \hat{T} , defined by

$$\hat{T}(n) := \max_{\omega \in \Omega_n} \mathbb{E}(T \mid \omega_0 = \omega),$$

The fact that the cooling process almost surely stops in a finite time now reads

$$\hat{T}(n) < \infty.$$

Conversely, since the volume of a tiling decreases at most by one at each step, the worst expected convergence time is bounded below by the maximal volume of tilings in Ω_n . With the hexagonal-shaped chosen domain, one checks that it yields the lower bound²

$$\hat{T}(n) = \Omega(n\sqrt{n}).$$

Our goal is to obtain a theoretical tight bound for the worst expected convergence time. Numerical experiments suggest that it is quadratic and correspond to tilings of maximal volume (Fig. 5). We thus conjecture:

$$\hat{T}(n) = \Theta(n^2).$$




Fig. 5. Left, markers show averages of the convergence time on numerical experiments, starting from the maximal volume tiling (thought to have the worst convergence time). We also drawn the curve cn^2 , $c = 6.689 \cdot 10^{-3}$, which matches rather well the markers. Right, dividing these experimental values by n^2 does not show a subpolynomial factor. We thus conjecture $\hat{T}(n) = \Theta(n^2)$.

² O , Ω and Θ are the usual Bachmann-Landau notations.

3 A theoretical upper bound

To obtain an upper theoretical bound on the expected convergence time, we will rely on the following probabilistic tool (proven, *e.g.*, in [4]):

Proposition 2. *Let $(x_t)_{t \geq 0}$ be a Markov chain over a space Ω . Assume that there is $\varepsilon > 0$ and a map $\phi : \Omega \rightarrow [a, b] \subset (0, \infty)$ such that, whenever $\phi(x_t) > a$:*

$$\mathbb{E}[\phi(x_{t+1}) - \phi(x_t) | x_t] \leq -\varepsilon.$$

Then, the expected value of the random variable $T := \min\{t \mid \phi(x_t) = a\}$ satisfies

$$\mathbb{E}(T) \leq \frac{b - a}{\varepsilon}.$$

We first need to introduce the notion of *triconvexity* (see Fig. 6):

Definition 1. *A dimer tiling is said to be triconvex if any segment of slope $0 \bmod \frac{2\pi}{3}$ which connects two vertices x and y contains only vertices whose height is between the heights of x and y . The triconvex hull $\bar{\omega}$ of a dimer tiling ω is the smallest triconvex dimer tiling greater than or equal to ω .*




Fig. 6. A non-triconvex tiling (left, with fault segments) and its triconvex hull (right).

Let us then associate with any map X defined over dimer tilings (for example, the volume V or the energy E) a map \bar{X} defined as follows:

$$\bar{X}(\omega) := X(\bar{\omega}).$$

This will help us to avoid ambiguities, especially when considering variations. Indeed, let us stress that $\Delta\bar{X}(\omega)$ and $\Delta X(\bar{\omega})$ are generally different, since

- $\Delta\bar{X}(\omega) = X(\bar{\omega}_a) - X(\bar{\omega})$, where ω_a is obtained by performing a flip on ω ;
- $\Delta X(\bar{\omega}) = X(\omega_b) - X(\bar{\omega})$, where ω_b is obtained by performing a flip on $\bar{\omega}$.

Now, in order to use Prop. 2, let us define a map ϕ over dimer tilings by:

$$\phi := 4V + E$$

We first consider the most simple case:

Lemma 1. *If ω is a triconvex dimer tiling with only one island, then*

$$\mathbb{E}[\Delta\phi(\omega)] \leq -\frac{12}{F(\omega)}.$$

Proof. Follow clockwise the boundary of the island: two consecutive edges make either a *salient* or a *reflex* angle, depending on whether we turn leftwards or rightwards (see Fig. 7). If there is a flip around some vertex x , then the six-triangle hexagon of center x has at least 3 error edges on its boundary. If these error edges are consecutive, then they form $i \geq 2$ consecutive similar angles, and performing the flip yields

- if the i angles are salient, then $\Delta V = -1$ and $\Delta E = 4 - 2i$, that is, $\Delta\phi = -2i$. This yields $\Delta\phi = -2$ per salient angle.
- if the i angles are reflex, then $\Delta V = 1$ and $\Delta E = 4 - 2i$, that is, $\Delta\phi = 8 - 2i$. This yields $\Delta\phi \leq +2$ per reflex angle.

If there are only such flips, then the claimed bound follows by checking by induction on its length that the boundary of an island (triconvex or not) has always six more salient than reflex angles. Otherwise, consider a flip in x such that the error edges on the boundary of the six-triangle hexagon of center x are non consecutive. At least two of these edges must be parallel. Thus, by triconvexity, the vertex x is in ω . Hence, such a flip decreases by 1 the volume (and disconnects the island). In particular, it decreases ϕ by at least 4, so that it can only improve the bound. \square




Fig. 7. A single triconvex island with 24 salient angles (downwards triangles) and 18 reflex angles (upwards triangles). A number in a circle shows the variation of ϕ if a flip is performed around this circle. There are $F(\omega) = 12$ performeable flips, and one computes $\mathbb{E}(\Delta\phi) = -\frac{3}{2} \leq -\frac{12}{F(\omega)}$. The bound is not tight because there is one flip (in the thick circle) which disconnects the island.

Since islands are not always triconvex, we need the following technical lemma, which relies the variations $\Delta\bar{\phi}(\omega)$ and $\Delta\phi(\bar{\omega})$ (see Fig. 8):

Lemma 2. *If ω is a dimer tiling with only one island, then*

$$F(\omega)\mathbb{E}[\Delta\bar{\phi}(\omega)] \leq F(\bar{\omega})\mathbb{E}[\Delta\phi(\bar{\omega})]$$

Proof. Let ΔX_y denotes the X -variation by performing a flip around y . Proving the lemma is then equivalent to prove

$$\sum_{x \in F(\omega)} \Delta\bar{\phi}_x(\omega) \leq \sum_{x \in F(\bar{\omega})} \Delta\phi_x(\bar{\omega}).$$

This easy follows from these three sublemmas:

- If $x \in F(\omega) \setminus F(\bar{\omega})$, then $\Delta\bar{\phi}_x(\omega) = 0$ since the flip does not modify $\bar{\omega}$.
- If $x \in F(\bar{\omega}) \setminus F(\omega)$, then $\Delta\phi_x(\bar{\omega}) > 0$.

We first prove that $\Delta V_x(\bar{\omega}) = 1$. This is clear if x is not in the island of $\bar{\omega}$. Otherwise, that is, if x is in the island of $\bar{\omega}$, then it is on a segment whose endpoints are in the island of ω and whose height are less or equal to the one of x . By performing the flip in x , these endpoints are unchanged. Hence, the triconvexity of $\bar{\omega}$ ensures that the height of x in $\bar{\omega}$ has increased, that is, $\Delta V_x(\bar{\omega}) = 1$.

We then prove that $\Delta E_x(\bar{\omega}) = 0$. If $\Delta E_x(\bar{\omega}) < 0$, then x is surrounded by at least 4 error edges, among which at least two are parallel. If x is not in the island of $\bar{\omega}$, this would contradict the triconvexity of $\bar{\omega}$. If x is in the island of $\bar{\omega}$, thus in the island of ω , then x would be in $F(\omega)$, what contradicts our initial assumption. Hence, $\Delta E_x(\bar{\omega}) = 0$ (since the cooling never increases the energy).

- If $x \in F(\omega) \cap F(\bar{\omega})$, then $\Delta\bar{\phi}_x(\omega) \leq \Delta\phi_x(\bar{\omega})$. Indeed, the only problem is that performing this flip could change much more the triconvex hull of ω than ω itself. Assume, first, that $\Delta E_x(\omega) = 0$. In this case, x is surrounded by exactly three consecutive error edges (on the boundary of the six-triangle hexagon modified by the flip). By performing the flip, we change these error edges by the three other ones. In particular, to each new error corresponds a parallel old error, so that the the triconvex hull after and before the flip differ only in x . Thus, $\Delta\bar{\phi}_x(\omega) = \Delta\phi_x(\bar{\omega})$.

Assume, now, that $\Delta E_x(\omega) < 0$. In this case, x is surrounded by $i \geq 4$ error edges. Since two of these edges are parallel, x is in the island of ω (otherwise it would not be in $F(\bar{\omega})$). One thus has $\Delta\phi_x(\bar{\omega}) = -4 - 2(i - 2) = -2i$. If x is still in the island of $\bar{\omega}$, then $\Delta\bar{\phi}_x(\omega) = 0$. Otherwise, once the flip performed, the new convex hull of the island has lost at least x , but maybe much more, so that we can just ensures that $\Delta\bar{\phi}_x(\omega) \leq -2i$. This shows $\Delta\bar{\phi}_x(\omega) \leq \Delta\phi_x(\bar{\omega})$. \square

In particular, this result combined with Lem. 1 applied to $\bar{\omega}$ yield:

$$\mathbb{E}[\Delta\bar{\phi}(\omega)] \leq \mathbb{E}[\Delta\phi(\bar{\omega})] \leq -\frac{12}{F(\omega)}.$$

Let us now extend this result to the case of several islands. The only problem which can occurs is that flips can merge islands. However, one big advantage of considering $\bar{\phi}$ instead of ϕ is that mergings becomes rather trivial, as we will see.




Fig. 8. An island of a tiling ω . Its triconvex hull (thick boundary) corresponds to Fig. 7. Flips are indicated by circles, triangles or squares, respectively depending whether they are in $F(\omega) \setminus F(\bar{\omega})$, $F(\bar{\omega}) \setminus F(\omega)$ or $F(\omega) \cap F(\bar{\omega})$. Lemmap 2 relies $\Delta\bar{\phi}(\omega)$ and $\Delta\phi(\bar{\omega})$.

Lemma 3. *If ω is a dimer tiling whose islands have all the same height, then*

$$\mathbb{E}[\Delta\bar{\phi}(\omega)] \leq -\frac{12}{F(\omega)}.$$

Proof. If there is no flip which can merge two or more islands of $\bar{\omega}$, then the result just follows from Lem. 1 and 2. Otherwise, consider such a flip. It is thus performed around a vertex, say x , which is not in $\bar{\omega}$. The triconvexity of $\bar{\omega}$ ensures that this is possible only if error and non-error edges alternate on the boundary of the six-triangle hexagon of center x . This flip thus merges exactly three islands of $\bar{\omega}$. Then, the triconvexity of $\bar{\omega}$ ensures that these islands are necessarily stick-shaped, that is, formed of aligned six-triangle hexagons (Fig. 9, left). At least the endpoints of these stick-shaped islands are in ω , and one easily checks that the bound still holds (Fig. 9, right). \square




Fig. 9. A flip which merges islands of $\bar{\omega}$ necessarily merges exactly three islands, which are moreover stick-shaped (left: $\bar{\omega}$). The flips in ω which modify $\bar{\omega}$ are such that ϕ decreases on expectation (right: ω).

We are now in a position to prove:

Theorem 1. *The worst expected convergence time satisfies:*

$$\hat{T}(n) = O(n^2\sqrt{n}).$$

Proof. Let us consider the highest level of the tiling (that is, the highest islands). The only problem to apply Lem. 3 is that some flips can be unperformable, due to lower islands. However, only flips which would increase the volume of $\bar{\omega}$ become unperformable. Moreover, these flips could not decrease the energy of $\bar{\omega}$ (we have already seen in the proof of Lem. 2 that a flip which increases the volume of a triconvex island cannot modify its energy). The bound of Lem. 3 thus holds *a fortiori* for this highest level. Since both the volume and the energy of the highest level of the tiling are at most n , and $F(\omega) \leq n$, one can apply Prop. 2 to $\bar{\phi}$ with $\varepsilon = 1/n$: this yields that the highest level of the tiling disappears in time $O(n^2)$. The claimed bound then follows since the tiling has $O(\sqrt{n})$ levels. \square

4 Prospects

Towards a tight bound

Recall that numerical experiments led us to conjecture that the expected convergence time is in $O(n^2)$. However, Th. 1 proves only a $O(n^2\sqrt{n})$ upper bound. This means that we lost a factor \sqrt{n} in the convergence time analysis. We see at least two possibly large approximation in our proof, informally described below.

First, we have considered triconvex hulls of islands instead of real islands. The main advantage is that we thus need to consider neither holes (hence we can suppose that an island has a simple boundary, as in Lem. 1) not complicated merging of islands (the merging analysed in Lem. 3 are very simple). The disadvantage is that, by considering triconvex hulls, we ignore the role of flips performed *inside* hulls (they just appear in the bound in the $F(\omega)$ term). However, numerical experiments suggest that the cooling of a tiling is not significantly faster than the cooling of its triconvex hull. Hence, it is not clear that we lost that much in considering triconvex hulls.

Second, in the proof of Th. 1, we applied Lem. 3 only the highest level of the tiling. This way, we make as if the $F(\omega)$ factor in the bound of Lem. 3 would count only flips on this level. But, of course, $F(\omega)$ counts all the flips of the tiling. In particular, one can reasonably expect that flips performed on lower levels act similarly, *i.e.*, that one could multiply the bound of Lem. 3 by the number of levels. Since they are around \sqrt{n} levels, this would greatly improve the bound. However, the problem is that a flip which could be performed on some lower island if they were not upped islands, can not be any more performable. This would not be a problem if such flips could only increase $\bar{\phi}$, but we have examples where the converse occurs. In other words, a lower island is forced to growth because of a higher island. We can maybe get round this problem, since the

higher island will rapidly shrink and no longer force the lower one to growth, but analyzing such a phenomenon seems to be very tedious...

Last, let us point out that the simple theoretical upper bound we provided is $\Omega(n\sqrt{n})$, that is, we also need to gain a factor \sqrt{n} .

Average expected convergence time

In this paper, we only consider the worst expected convergence time, $\hat{T}(n)$. However, it is also natural (even more, according to the physical motivations exposed in the introduction), to consider the average expected convergence time:

$$\bar{T}(n) := \frac{1}{\#\Omega_n} \sum_{\omega \in \Omega_n} \mathbb{E}(T \mid \omega_0 = \omega).$$

In this case, numerical experiments (Fig. 10) led us to conjecture:

$$\bar{T}(n) = \Theta(n\sqrt{n}).$$




Fig. 10. Analog of Fig. 5 for the average expected convergence time. Here, the drawn curve (left) has equation $cn\sqrt{n}$, $c = 1.669 \cdot 10^{-2}$. We thus conjecture $\bar{T}(n) = \Theta(n\sqrt{n})$.

In order to bound the average expected convergence time by similar technics as for the worst expected convergence time, we first need to compute the average values of some functions, *e.g.*, the volume V , the energy E or the number of levels H . This is what we did in the one-dimensional case in [2]. However, this two-dimensional case is much harder, in the spirits of works by Propp & al [6]. Let us just mention, to conclude, that numerical experiments indicate that both the average volume and the average energy of a tiling are linear (in the number of tiles), while the number of levels is logarithmic!

Acknowledgments. We would like to thank Olivier Bodini (LIP6, Paris), Éric Rémila (LIP, Lyon) and Mathieu Sablik (LATP, Marseille) for useful discussions.

References

1. O. Bodini, Th. Fernique, D. Regnault, *Crystallization by stochastic flips*, to appear in Journal of Physics: Conference Series.
2. O. Bodini, Th. Fernique, D. Regnault, *Stochastic flips on two-letter words*, Analytic Algorithmics & Combinatorics (2010), p. 48–55.
3. O. Bodini, Th. Fernique, É. Rémila, *A Characterization of flip-accessibility for rhombus tilings of the whole plane*, Inf. Comput. **206** (2008), p. 1065–1073.
4. N. Fatès, M. Morvan, N. Schabanel, É. Thierry, *Fully asynchronous behavior of double-quiescent elementary cellular automata*, Theor. Comput. Sci. **362** (2006), pp. 1–16.
5. C. Janot, *Quasicrystals: The State of the art*, Oxford University Press, 1992.
6. H. Cohn, M. Larsen, J. Propp, *The Shape of a typical boxed plane partition*, New-York J. Math. **4** (1998), p. 137–166.
7. W. P. Thurston, *Conway tilings group*, Amer. Math. Month. **97** (1990), p. 757–773.